

CLAIMS

What is claimed is:

- 1 1. A method for responding to a request to transfer data from a virtual
2 computer system to a computer network, the virtual computer system comprising a
3 plurality of NICs, the method comprising:
4 determining NIC management information related to possible NICs over
5 which the data may be transferred;
6 determining VM-specific information related to one or more virtual
7 machines (VMs) in the virtual computer system;
8 based on the NIC management information and the VM-specific
9 information, deciding whether to transfer the data; and
10 if a decision is made to transfer the data, then, based on the NIC
11 management information and the VM-specific information, selecting a NIC over
12 which to transfer the data.
- 1 2. The method of claim 1, in which the VM-specific information indicates an
2 amount of network bandwidth that is allocated to a VM that requested the data transfer.
- 1 3. The method of claim 2, in which a decision is made not to transfer the data
2 because transferring the data would cause the VM's allocation of network bandwidth to
3 be exceeded.
- 1 4. The method of claim 1, in which the VM-specific information indicates the
2 priority of the VM that requested the data transfer relative to the priorities of other virtual
3 machines.
- 1 5. The method of claim 1, in which the NIC management information
2 indicates the available NICs over which the data may be transferred.

1 6. The method of claim 5, in which the NIC management information further
2 indicates a pending data transfer load for each of the available NICs over which the
3 data may be transferred.

1 7. The method of claim 1, in which a load distribution function, based on the
2 NIC management information and the VM-specific information, is used in selecting a
3 NIC over which to transfer the data.

1 8. The method of claim 7, in which a first VM's data transfer requests are
2 substantially always routed over a first NIC as long as the first NIC is available, and a
3 second VM's data transfer requests are substantially always routed over a second NIC
4 as long as the second NIC is available, but the first VM's data transfer requests are
5 routed over the second NIC if the first NIC is not available, and the second VM's data
6 transfer requests are routed over the first NIC if the second NIC is not available.

1 9. The method of claim 8, in which the first VM's data transfer requests are
2 distinguished from the second VM's data transfer requests by reference to a source
3 physical address contained in a header of each data transfer request.

1 10. The method of claim 1, in which the NIC management information
2 indicates whether a failover is occurring on one of the NICs over which the data could
3 otherwise have been transferred.

1 11. The method of claim 10, in which the VM that has requested the data
2 transfer is temporarily suspended if a failover is occurring on one of the NICs over which
3 the data could otherwise have been transferred.

1 12. The method of claim 1, wherein, if a decision is made not to transfer the
2 data, the data is placed on a queue for transferring at a later time.

1 13. The method of claim 1, wherein, if a decision is made not to transfer the
2 data, a further decision is made whether to suspend the VM that requested the data
3 transfer.

1 14. The method of claim 1, wherein, if a decision is made not to transfer the
2 data, a further decision is made whether to migrate the VM that requested the data
3 transfer to another computer system.

1 15. A computer program embodied in a tangible medium, the computer
2 program executing in a virtual computer system in support of a plurality of VMs, the
3 virtual computer system having access to a plurality of physical NICs connected to a
4 computer network, the computer program comprising:
5 a resource manager for sharing system resources between the plurality of
6 VMs; and
7 a NIC manager for routing data from the plurality of VMs through the
8 plurality of NICs to the computer network,
9 wherein the computer program determines VM-specific information and
10 NIC management information and the NIC manager makes NIC management
11 decisions based on both the VM-specific information and the NIC management
12 information.

1 16. The computer program of claim 15, wherein the VM-specific information
2 and the NIC management information are used together when making decisions
3 regarding the sharing of system resources.

1 17. The computer program of claim 16, wherein the decisions regarding the
2 sharing of system resources are based on a network bandwidth shaping filter approach.

1 18. The computer program of claim 16, wherein the decisions regarding the
2 sharing of system resources are based on a proportional-share approach.

1 19. The computer program of claim 15, further comprising a VM manager for
2 controlling the general operation of the plurality of VMs.

1 20. The computer program of claim 19, wherein the VM-specific information
2 and the NIC management information are used together when making decisions
3 regarding the management of the plurality of VMs.

1 21. The computer program of claim 19, wherein the VM manager and the
2 resource manager are both implemented in a single software unit.

1 22. The computer program of claim 19, wherein the VM manager, the
2 resource manager and the NIC manager are integrated together within a kernel.

1 23. The computer program of claim 19, further comprising a device emulator
2 for presenting a virtual NIC to each of one or more of the VMs to provide access to the
3 plurality of physical NICs.

1 24. A method for selecting a NIC to be used for a data transfer between a VM
2 within a virtual computer system and a computer network, the virtual computer system
3 comprising a plurality of NICs and a plurality of VMs, the method comprising:
4 determining which VM within the virtual computer system is involved in the
5 data transfer;
6 determining a plurality of available NICs over which the data may be
7 routed; and
8 based on the particular VM involved in the data transfer and the available
9 NICs over which the data may be routed, selecting a NIC over which to route the
10 data.

1 25. The method of claim 24, wherein the selection of a NIC over which to
2 route the data is further based on information regarding pending data loads on the
3 available NICs over which the data may be routed.

1 26. The method of claim 24, wherein the selection of a NIC over which to
2 route the data is further based on relative priorities for the plurality of VMs.

1 27. The method of claim 24, wherein the selection of a NIC over which to
2 route the data is further based on system resource allocations for the plurality of VMs.

1 28. The method of claim 24, wherein the selection of a NIC over which to
2 route the data is based on a load distribution algorithm.

1 29. The method of claim 28, wherein the load distribution algorithm is a load-
2 balancing algorithm.

1 30. A method for routing data between a virtual computer system and a
2 computer network, the virtual computer system comprising a first VM and a second VM,
3 the virtual computer system also comprising a first NIC and a second NIC for
4 connecting to the computer network, the method comprising:

5 for each data transfer request:

6 determining which VM within the virtual computer system is
7 involved in the requested data transfer; and

8 if the first VM is involved in the requested data transfer, routing the
9 data over the first NIC; or

10 if the second VM is involved in the requested data transfer, routing
11 the data over the second NIC.

1 31. A method for transferring data from a VM in a virtual computer system to a
2 computer network, the virtual computer system comprising a first physical NIC and a
3 second physical NIC for connecting to the computer network, the method comprising:

4 presenting a virtual NIC to the VM through which the VM may attempt to
5 transfer data to the network;

6 receiving a first set of data and a second set of data, both of which were

7 sent by the VM to the virtual NIC for transference to the network; and
8 based on a load distribution algorithm, transferring the first set of data to
9 the network through the first physical NIC and transferring the second set of data
10 to the network through the second physical NIC.

1 32. The method of claim 31, wherein the load distribution algorithm is a load-
2 balancing algorithm.

1 33. The method of claim 32, wherein the load-balancing algorithm implements
2 a round-robin function.

1 34. The method of claim 32, wherein the load-balancing algorithm is based on
2 a pending data transfer load for each of the first and second physical NICs.

1 35. The method of claim 31, wherein the load distribution algorithm is based
2 on NIC management information.

1 36. The method of claim 31, wherein the load distribution algorithm is based
2 on VM-specific information.

1 37. The method of claim 31, wherein the load distribution algorithm is based
2 on both NIC management information and VM-specific information.

1 38. A method for transferring data from a VM in a virtual computer system to a
2 computer network, the virtual computer system comprising a first physical NIC and a
3 second physical NIC for connecting to the computer network, the method comprising:
4 presenting a virtual NIC to the VM through which the VM may attempt to
5 transfer data to the network;
6 receiving a plurality of data, which were sent by the VM to the virtual NIC
7 for transference to the network; and
8 transferring each of the plurality of data to the network through either the

first physical NIC or the second physical NIC, wherein if the first physical NIC is not available for use, the data are transferred to the network over the second physical NIC, and if the second physical NIC is not available for use, the data are transferred to the network over the first physical NIC.

39. The method of claim 38, wherein if both the first physical NIC is available for use and the second physical NIC is available for use, the data are transferred to the network using a physical NIC that is determined according to a load distribution algorithm.

40. A method for managing network access between a virtual computer system and a computer network, the virtual computer system comprising a plurality of VMs, the virtual computer system being susceptible to possible adverse effects from a Denial of Service attack, the method comprising:

- monitoring data that is transferred between the virtual computer system and the computer network for an indication of a possible Denial of Service attack; and
- if an indication of a possible Denial of Service attack is detected, suspending one or more of the VMs, to reduce the risk of adverse effects on one or more other VMs.

41. A method for managing network access between a VM and a computer network, the VM executing in a first physical computer system having access to a first connection to the computer network, the first connection being susceptible to possible adverse effects from a Denial of Service attack, the method comprising:

- monitoring data that is transferred between the VM and the computer network for an indication of a possible Denial of Service attack; and
- if an indication of a possible Denial of Service attack is detected, migrating the VM to a second physical computer system having access to a second connection to the computer network.

42. The method of claim 41, wherein the step of monitoring the data involves monitoring the quantity of network traffic that is transferred between the VM and the computer network.

43. The method of claim 41, wherein the first and second physical computer systems are connected directly to the computer network.

44. A method for managing network access between a virtual computer system and a computer network, the virtual computer system comprising a VM, the virtual computer system being susceptible to possible adverse effects from a Denial of Service attack, the method comprising:
monitoring data that is transferred from the computer network to the virtual computer system for an indication of a possible Denial of Service attack; and
upon detecting data directed toward the VM, where the data indicates a possible Denial of Service attack, delaying the processing of the data by the VM to allow for one or more defensive measures against the possible Denial of Service attack.

45. The method of claim 44, wherein the processing of the data by the VM is delayed by delaying an interruption to the VM regarding receipt of the data frame.

46. The method of claim 44, wherein the processing of the data by the VM is delayed by temporarily suspending the VM.

47. The method of claim 44, wherein the processing of the data by the VM is delayed by delaying scheduling of the VM for CPU execution time.

1 48. The method of claim 44, wherein the one or more defensive measures
2 comprises failing over from the use of a first set of one or more NICs to the use of a
3 second set of one or more NICs.